National Security Applications Co-Design:
A Framework for an ASC Tri-lab Project

Introduction

The challenges facing Advanced Simulation and Computing (ASC) *Engineering and Physics Integrated Codes* (EPICs, or simply "integrated codes") designers and developers and platform acquirers during this decade—and into the next—are daunting. It has been well understood for several years that a paradigm shift in computer design is under way– one requiring developers to treat memory as scarce, locality as critical, and compute power and parallelism as a relatively cheap resource. In addition, electrical power requirements for the next 1000-fold leap in computing power are such that vendors are proposing exotic and creative ways to extract an increasing number of flops-per-watt. Whether through the use of accelerators such as GPUs, or massive parallelism using shared memory, coming architectural changes will necessitate significant redesign, and quite possibly complete rewrites, of many EPICs. ASC EPICs are somewhat unique in the DOE due to their extreme size, complexity, and programmatic mission, making it impractical to pause and rewrite upon delivery of each new platform in the next generation of architectures.

Given the now well-appreciated challenges of achieving efficient operations on next generation platforms, application code developers cannot simply specify a set of requirements for platform providers to meet and then divorce themselves from the platform acquisition process. Similarly, the applications code teams cannot quietly accept whatever platform is provided by vendor partners without some input into the tradeoff decisions made that affect programming models, system software, and performance. In this white paper, we use the term "National Security Applications Co-Design Project" (or *NSApp CDP*) to refer to the collection of three NNSA/ASC Co-Design Hubs (CDHubs) under a standard operations framework. The availability of funding may elevate either this ASC project or each laboratory hub to *Co-design Center* status such as those in the Office of Science ASCR program, but for now we use alternate terminology to make clear that this is a nascent effort on the part of ASC to bootstrap an effort using limited existing resources.

The NSApp CDP provides the vehicle to engage in the collaborative process addressing the challenges outlined above, focusing on the unique aspects of integrated codes at the NNSA labs. We propose a high-level structure for a tri-lab (LLNL, LANL, SNL) National Security Applications Co-Design Project that focuses on the following key needs from the ASC EPIC mission:

- Next generation platforms delivered by the ASC Program must allow effective exploitation by ASC EPICs, not just science codes, to meet mission requirements.
- The emphasis on physics coupling in the ASC integrated codes presents an additional challenge that will likely be addressed only by the NSApp CDP.

- Large codes often stress the software stack much more than smaller codes, whether it's tools (debuggers, hardware simulators, etc.), compilers, or operating systems.
- Due to critical mission requirements, at-scale resiliency of the entire system – consisting of the hardware, the software stack, and the application codes – must be addressed so that simulations of interest can be completed when needed.
- Because of the size of these ASC EPICs (often > 1M lines), and the necessity to maintain the same code base for multiple platforms, code reuse and performance portability is a high priority.

At first we expect each laboratory hub to set its own direction in the effective use of proprietary, non-standard, or vendor-specific methods. Eventually a standard programming model should emerge with the goal of making hardware-specific coding transparent to all application programmers.

The future is littered with trade-offs, and give and take will be needed from both the hardware and software developers. Understanding and influencing these trade-offs is a principal role of the NSApp CDP.

Why the need for an ASC focused Co-Design Project

National security integrated codes uniquely need to be in service for many years, often a decade or more, because of their complexity and long development time. Many codes now reaching maturity in the design community were started 15+ years ago with the beginning of the ASCI program. The codes must model a variety of physical configurations and allow the end user flexibility in their operation, sometimes through the use of code steering. And unlike many science codes, the users of EPICs are typically not code developers. They demand a level of robustness that necessarily requires rigor in software quality assurance practices, and careful attention to usability factors affecting a community of users, such as reproducibility, verification and validation, and ease of use.

Typically, ASC EPICs couple together more physics packages than most application codes targeted by the ASCR Co-design Centers (CDCs). It is likely, for simplicity's sake, that NSApp CDP's interaction with vendor partners – the co-design process itself – will center on individual algorithms and technologies or, perhaps two physics packages working in concert. The increased emphasis on physics coupling in the ASC integrated codes presents an additional challenge that will likely be addressed only by the NSApp CDP. Likewise, when a synergy is not possible due to a unique characteristic of NNSA's codes, NSApp CDP will pursue co-design research independently with the appropriate vendor partner, with the goal of sharing generalized results (i.e., those that do not compromise vendor intellectual property) with the DOE co-design community whenever possible.

The various exascale workshops recently held by both ASCR and ASC have produced requirements for next generation platform providers from various groups of stakeholders. A central objective of the NSApp CDP will be refining these requirements to their essence while working with vendor partners and ensuring the resulting product will meet NNSA mission requirements. From the perspective of an applications code developer, an example of one such requirement is reproducibility to enable debugging. A vendor may offer an alternative solution that still enables effective debugging even without strict reproducibility. It is trade-offs of this type that must be explored in the co-design process.

An additional metric critical to ASC success is performance portability. These applications must operate efficiently on a variety of platforms to meet programmatic demands, and we cannot afford to maintain multiple versions of our large (1-4M LOC) applications tuned for different architectures. In the short term, such coding will probably be necessary, so it must be minimized and isolated using good software engineering practices, and explored in conjunction with the community using proxy applications (see below). Another important requirement is the development of a standard programming model, akin to MPI, that isolates hardware specific coding from most application programmers.

However, due to classification concerns, details of nuclear weapons simulation codes are protected. Other codes of interest to national security, however, share some of the important characteristics and algorithms present in these types of weapons codes. These algorithms and technologies – and their applicability to the next generation platforms – will be explored in the NSApp CDP.

The ASC Program benefits from each of the three NNSA Laboratories proposing a unique co-design hub, with each lab working closely with the vendor partner(s) of their choosing – likely based on procurements of advanced architecture systems. While the mission space of LANL and LLNL are nearly identical, diversity in approaches will reduce technical risk associated with meeting the challenge of executing the ASC mission on next generation platforms. Likewise, the SNL science and engineering codes fill a unique role in the NNSA mission, and an independent co-design project at SNL will help to ensure that their particular issues are satisfied. If three NSApp co-design hubs are pursued, it is crucial that all three communicate lessons learned and concerns in a timely manner because, once platforms are delivered, LLNL codes will be expected to run effectively on ACES platforms, and, conversely, LANL and SNL codes must run on LLNL platforms.

As is frequently the case when dealing with any new technological challenge, multiple vendors are employing multiple solutions to manage machine power consumption, and it is not clear which strategy will eventually achieve widespread adoption throughout the industry. Therefore, mitigating risk by establishing multiple hubs that can deal with multiple vendors is an important component to the CDC strategy. Communication between the hubs is also critical to ensure that programming strategies and standards adopted by the labs will be successful on the platforms likely to be delivered through the co-design collaborations.

The NSApp CDP will thus focus on the technologies required by EPICs. The simulation codes used to model the behavior of nuclear weapons are a type of integrated code, and are the ultimate product of the ASC Program. *Next generation platforms delivered by the ASC Program must allow effective exploitation by ASC's nuclear weapons simulation codes.*

Science Codes and Integrated Codes

Codes of interest to the NNSA National Security community generally fall into two broad categories: *Science Codes* and *Engineering and Physics Integrated Codes (EPICs).* Science codes usually delve into a single type of physics or phenomena, performing simulations in regimes that would be either impossible or impractical to capture experimentally in order to develop physical data that, in the end, are used by the integrated codes. In contrast, integrated codes typically combine multiple physics to represent complex behaviors of physical systems of interest.

Examples of science codes include molecular and dislocation dynamics codes that characterize materials in extreme conditions, atomic physics codes that describe the interaction of material with radiation, and nuclear physics codes that produce nuclear interaction cross sections.

Science codes are critical in NNSA's pursuit of a predictive simulation capability. They provide insight through direct numerical simulation of subscale mechanisms. The crucial constitutive models required to solve the hydrodynamics in integrated codes are produced by science codes. Difficulties arise when the material constitution changes on the same time-scale as the hydrodynamics, such as during a phase transformation. Handling these challenges may necessitate closer coupling of science and design codes (perhaps through dynamic spawning of a science code from the EPIC) that will require significantly more computing resources. The ASCR CDCs that focus on simulation technologies related to science codes, for example LANL/LLNL's *Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx)*, are of particular interest to ASC and its Physics and Engineering Models component.

Partnerships and Governance

The NSApp CDP will be coordinated out of NNSA ASC HQ, with co-leadership from the individual co-design hubs at the three NNSA labs. The individual co-design hubs will also be coordinating with the vendor community as well as the ASCR Co-design Centers. In order to best address issues such as the development of common legal frameworks, and handling of vendor and DOE intellectual property, convening of workshops, and funding, we propose to work with ASCR on a lightweight governance model across DOE for these interactions.

Coordination with ASCR Co-Design Centers

The three ASCR CDCs awarded in June 2011 are:

- ANL's *Center for Exascale Simulation of Advance Reactors (CESAR)*
- LANL/LLNL's *Exascale Center for Materials in Extreme Environments (ExMatEx)*
- SNL/LBNL's *Center for Exascale Simulation of Combustion in Turbulence (ExaCT)*

Many of the algorithms and technologies to be addressed by the NSApp CDP are also shared with subject areas being pursued by ASCR CDCs. For example, SNL/LBNL's *Center for Exascale Simulation of Combustion in Turbulence (ExaCT)* holds interest because its treatment of disparate timescales is similar to that necessary in modeling high explosives.

NSApp CDP will partner closely with these ASCR CDCs to ensure that crosscutting algorithms and technologies are shared to avoid duplication of effort and to ensure the needs of NNSA are met. Examples of common crosscutting solutions include programming models and languages, I/O, operating systems, and portability.

In addition to the crosscuts, there are also application areas common to the NSApp CDP and the ASCR CDCs. For example, *ExMatEx* is creating a task-based parallel proxy application to enable adaptive physics refinement within a hydrodynamics simulation code. Examples exist for all three ASCR CDCs. Coupling between the NSApp CDP and the ASCR CDCs will present to the vendor community the broader application workload expected for next generation platforms and will enable essential cross-fertilization of lessons learned between science and integrated codes.

DOE Coordination of Vendor Partnerships

NSApp CDP, like many of the ASCR CDCs, will conduct its investigation with its vendor partners through a tightly coupled cycle of application, algorithm and system software development, performance modeling, hardware simulation, and hardware design. Through the co-design process optimal design tradeoffs will be identified for hardware, the run-time environment, parallel file systems, physics algorithms and packages, and programming models. In early discussions between ASC, ASCR, and the vendor community, several issues have arisen that must be addressed in order for the co-design process to work effectively as envisioned.

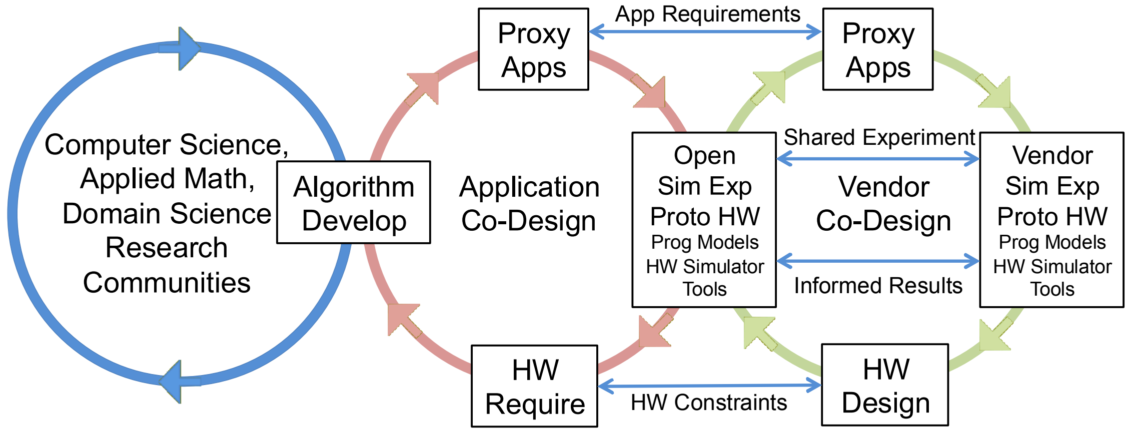

Figure 1 Notional Description of Co-design (courtesy: ASCR)

First, there must be a methodology in place to prevent contamination of vendor intellectual property (IP). Early discussions with vendors have raised a number of concerns regarding the potential release of highly sensitive non-disclosure information as well as learning of their competitors' trade secrets through inadvertent exposure. Several proposed solutions have been suggested by not yet agreed upon. These include "firewalling" individuals in the CDCs who are given access to "deep IP", something that CDCs are reluctant to accept. Likewise the abstraction of efforts through architectural simulators and use of proxy apps have been suggested, but may not provide the intended impact. The legal and technical complexities of this issue alone demands high-level coordination across the DOE co-design centers.

Second, there must be a way to manage the *MxN* interactions between *M* CDCs and *N* vendor teams. This coordination will benefit the CDCs by lowering the barriers for initial vendor engagement by keeping points-of-contact, existing CDC-vendor relationships, and ensuring maximum leverage across the community.

The ASCR CDCs have proposed the development of an *Exascale Co-Design Consortium (ECDC)* to study and support these two critical requirements as well as any additional issues between DOE and the vendor communities that arise in the future. **The NSApp CDP will agree to participate in the ECDC as a first step in systematizing the CDC approach across the DOE.**

A High Level Model for DOE Co-Design Coordination

The suggested hierarchical model is for a coordinated DOE Co-Design Effort that identifies and exploits commonalities and identifies gaps across *all* co-design hubs and centers in NNSA/ASC and SC/ASCR, and. The highest-level effort will mainly consist of program managers from NNSA/ASC and SC/ASCR coordinating activities, and is intended to be lightweight – focused on ensuring communication between the various hubs and centers maximally leverages the strengths of each, and minimizes duplicate effort.

The following diagram is a high level overview of a proposed structure for managing co-design hubs in NNSA with the centers across DOE. The goal is for NNSA to maximally leverage co-design work going on in ASCR, which ensuring NNSA can pursue its own unique requirements under a coordinated effort.

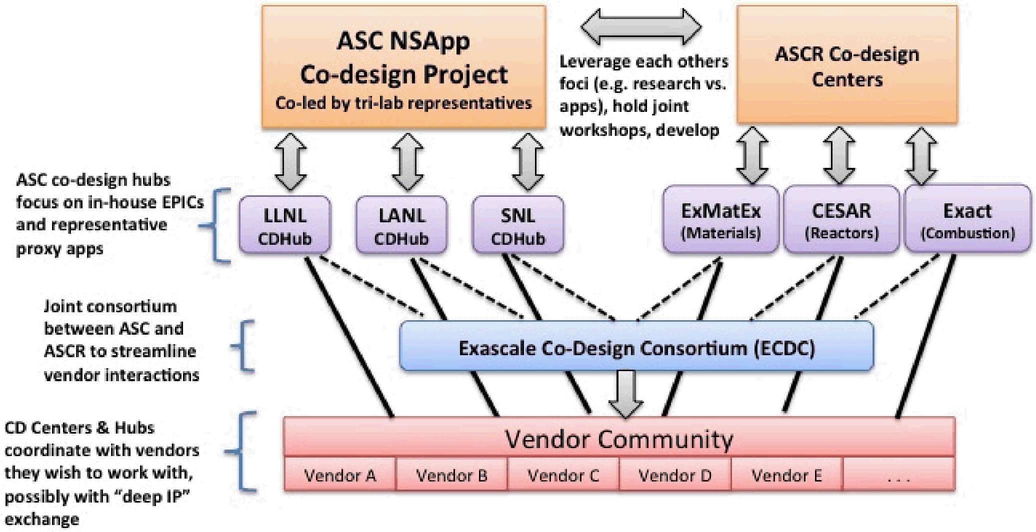

Figure 2 - Governance and coordination between ASC, ASCR, and vendor community

The NNSA CDP will be represented by three hubs at each NNSA laboratory. These individual co-design hubs will represent the vehicles for deliverables, and interactions with the vendor community. Each NNSA co-design hub will ultimately be charged with enabling their own lab's success in deploying a next generation set of integrated codes, with guidance and oversight from the NSApp CDP. The ECDC will act as the conduit to efficiently manage relationships between DOE efforts and vendor partners, as well as ensure strong continued coordination between ASC hubs and the ASCR CDCs.

It is anticipated that, at least initially, most of the co-design efforts in NSApp CDP will focus on development of benchmark applications of the type described in section "**Error! Not a valid bookmark self-reference.**Proxy Applications". For the most part, we anticipate that many of these proxy apps can be shared across the NSApp CDP (i.e., with multiple vendors).

Proxy Applications

NSApp CDP will use representative proxy applications as the means to interact with our vendor partner(s) during the co-design process. Whenever possible, proxy apps are fully open and released without restriction, although occasionally certain algorithms of interest may fall under export-control restriction. Proxy applications can be grouped into categories of increasing sophistication and fidelity in relation to the actual applications (or packages) used in integrated codes. One possible taxonomy is:

1. **Kernels**: These are one or more small code fragments or data layouts that are used extensively by the applications and are deemed essential to perform optimally on exascale platforms. These are useful for testing programming methods and performance at the device level, and typically do not involve network communication (MPI).

2. **Skeleton apps**: These apps reproduce the data flow or communication patterns of a physics application or package, and make little or no attempt to investigate numerical performance. They are primarily useful in investigating network performance characteristics at large scale.
3. **Mini apps**: These apps combine contain some of the dominant numerical kernels (or subsets thereof) contained in an actual application and produce simplifications of physical phenomena. This category may also include libraries wrapped in a test driver providing representative inputs.
4. **Compact apps**: These apps are representative of the actual application. In some cases compact apps may be full-fledged physics packages, and as such are often restricted in their distribution.

Each lab has at its disposal a number of proxy applications that have been produced over the years as part of the ASC machine procurement process. These applications capture the essence of many important physics algorithms, but fall short of completely capturing the breadth of algorithms across NNSA EPICs. As such, an initial effort on gap analysis and identification of new proxy applications needed to envelope NNSA EPICs will be undertaken. The NSApp CDP will then compare this collection of algorithms, solvers, and applications to the scope of the ASCR CDCs to identify items that are unique to the NSApp CDP.

These applications will be used both by the vendors to understand the effects of hardware tradeoffs, but also by integrated code team members and DOE researchers wishing to explore and develop new technologies, runtime systems, languages, programming models, algorithms, tools, file systems, visualization techniques, etc. For example, proxy apps are a valuable way to expose new staff to the concepts of scientific computing and algorithms of importance to ASC in an open and manageable environment.

NSApp CDP Use of Existing Testbeds

The labs are currently deploying a number of testbeds that are representative of new architectures. It would make sense to develop efficient implementations of our miniapps using a variety of programming models and identify gaps. We need to share experiences related to both the architectures and the programming models. Providing the ASC co-design hubs access to the various testbeds would be an efficient and impactful step towards developing applications that can run on multiple exascale platforms in the future.

NSApp CDP Focus Areas

Physics and numerical algorithms that the NSApp CDP will investigate (and are candidates for algorithmic collaboration with ASCR CDCs) include

- Hydrodynamics
- Deterministic transport
- Monte Carlo transport
- Diffusion processes
- Coupled hypersonic aero-thermal
- Structural mechanics, shock and vibration, and failure models
- Thermal effects and thermal-structural interactions
- Electromagnetics

- Linear and non-linear solvers
- Subscale models

ASC integrated codes couple together many physics packages. Investigation of focus areas above will emphasize investigation in the context of multi-physics integration.

Next generation architectures will pose some stringent challenges to national security applications and the productivity of their code developers. Technologies the NSApp CDP will investigate include:

- Programming models and languages
- Runtime systems for orchestration of multi-physics computation
- Billion-way parallelism
- Accelerators and hybrid systems
- Input/output and file systems
- Resiliency and fault tolerance
- Data layout and movement (algorithmic studies)
- Non-SPMD (for example, MPMD-like models)
- Complex memory hierarchies
- Power limitations, and programmer-controlled energy use
- Memory limitations and shared memory
- Operating System trade-offs

This is also an opportunity to reexamine non-conventional algorithms that could be enabled by the new technology. Higher order schemes where more computation is performed for a given amount of data motion may be able to take advantage of the new technology. For example, high order finite element hydro schemes, despite their ability to achieve greater accuracy for a given number of zones, have historically not been worth the additional computational cost. Rather, running at finer resolution with low order schemes has provided faster turnaround for a given level of accuracy. That paradigm may change with architectures where flops are very cheap relative to memory access.

The CDP will also be used to identify and prioritize changes that can be implemented in future hardware capabilities. CDP investigations will be performed in a diverse set of technical areas: Architecture-aware algorithms, programming models, system software, hardware architectures, resiliency, power management, etc. The key tools we apply for these investigations are proxy applications, architectural simulations, and experimental and advanced architecture testbed platforms.

Conclusion

The ASC program must find a way to ramp up engagements with the vendor community in the co-design process. Co-design is already a fundamental attribute of ASC code development, and we should formalize this as soon as possible in order to collaborate effectively with ASCR co-design activities. Additional funding in the future could extend the scope of co-design efforts to allow for deeper interactions, enable more coordination with ASCR, and provide much needed "free energy" in enable the integrated code teams to explore the fast moving bow wave of changes sweeping through the code development community. Whether or not an official exascale program emerges within the next decade, it is clear that even commodity computing platforms arriving in the next several years will require dramatic changes in the way codes are developed. Fine-grained parallelism, resilience, SIMD/vectorization, multi-level memory models – these are just a few of the topics that the ASC engineering and physics integrated code teams must embrace as a new reality and address as part of their future.